3

Time Series Forecasting Techniques

Back in the 1970s, we were working with a company in the major home appliance industry. In an interview, the person in charge of quantitative forecasting for refrigerators explained that their forecast was based on one time series technique. (It turned out to be the exponential smoothing with trend and seasonality technique that is discussed later in this chapter.) This technique requires the user to specify three "smoothing constants" called α, β, and γ (we will explain what these are later in the chapter). The selection of these values, which must be somewhere between 0 and 1 for each constant, can have a profound effect upon the accuracy of the forecast.

As we talked with this forecast analyst, he explained that he had chosen the values of 0.1 for α, 0.2 for β, and 0.3 for γ. Being fairly new to the world of sales forecasting, we envisioned some sophisticated sensitivity analysis that this analyst had gone through to find the right combination of the values for the three smoothing constants to accurately forecast refrigerator demand.

However, he explained to us that in every article he read about this technique, the three smoothing constants were always referred to as α, β, and γ, in that order. He finally realized that this was because they are the 1st, 2nd, and 3rd letters in the Greek alphabet. Once he realized that, he "simply *took 1, 2, and 3, put a decimal point in front of each, and there were my smoothing constants."*

After thinking about it for a minute, he rather sheepishly said, "You know, it doesn't work worth a darn, though."

❖ INTRODUCTION

We hope that over the years we have come a long way from this type of time series forecasting. First, it is not realistic to expect that each product in a line like refrigerators would be accurately forecast by the same time series technique—we probably need to select a different time series technique for each product. Second, there are better ways to select smoothing constants than our friend used in the previous example. To understand how to better accomplish both of these, the purpose of this chapter is to provide an overview of the many techniques that are available in the general category of time series analysis. This overview should provide the reader with an understanding of how each technique works and where it should and should not be used.

Time series techniques all have the common characteristic that they are endogenous techniques. This means a time series technique looks at only the patterns of the history of actual sales (or the series of sales through time—thus, the term time series). If these patterns can be identified and projected into the future, then we have our forecast. Therefore, this rather esoteric term of endogenous means time series techniques look inside (that is, endo) the actual series of demand through time to find the underlying patterns of sales. This is in contrast to regression analysis, which is an exogenous technique that we will discuss in Chapter 4. Exogenous means that regression analysis examines factors external (or exo) to the actual sales pattern to look for a relationship between these external factors (like price changes) and sales patterns.

If time series techniques only look at the patterns that are part of the actual history of sales (that is, are endogenous to the sales history), then what are these patterns? The answer is that no matter what time series technique we are talking about, they all examine one or more of only four basic time series patterns: level, trend, seasonality, and noise. Figure 3.1 illustrates these four patterns broken out of a monthly time series of sales for a particular refrigerator model. The *level* is a horizontal sales history, or what the sales pattern would be if there were no trend, seasonality, or noise. For a product that is sold to a manufacturing concern as a component in another product whose demand is stable, the sales pattern for this product would be essentially level, with no trend, seasonality, or noise. In our example in Figure 3.1, however, the level is simply the starting point for the time series (the horizontal line), with the trend, seasonality, and noise added to it.

Trend is a continuing pattern of a sales increase or decrease, and that pattern can be a straight line or a curve.

Of course, any business person wants a positive trend that is increasing at an increasing rate, but this is not always the case. If sales are decreasing (either at a constant rate, an increasing rate, or a decreasing rate), we need to know this for forecasting purposes. In our example in Figure 3.1, trend is expressed as a straight line going up from the level.

Seasonality is a repeating pattern of sales increases and decreases that occurs within a one-year period or less ("seasonal patterns" of longer than one year are typically referred to as "cycles," but can be forecast using the same time series techniques). Examples of seasonality

Figure 3.1 Time Series Components

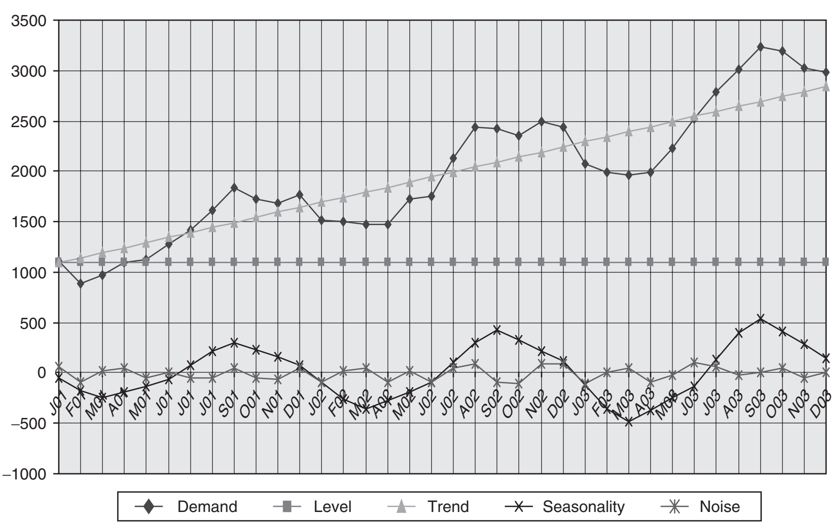

are high sales every summer for air conditioners, high sales of agricultural chemicals in the spring, and high sales of toys in the fall. The point is that the pattern of high sales in certain periods of the year and low sales in other periods repeats itself every year. When broken out of the time series in Figure 3.1, the seasonality line can be seen as a regular pattern of sales increases and decreases around the zero line at the bottom of the graph.

Noise is random fluctuation—that part of the sales history that time series techniques cannot explain. This does not mean the fluctuation could not be explained by regression analysis or some qualitative technique; it means the pattern has not happened consistently in the past, so the time series technique cannot pick it up and forecast it. In fact, one test of how well we are doing at forecasting with time series is whether the noise pattern looks random. If it does not have a random pattern like the one in Figure 3.1, it means there are still trend and/or seasonal patterns in the time series that we have not yet identified.

We can group all time series techniques into two broad categories—*open-model time series techniques* and *fixed model time series techniques*—based on how the technique tries to identify and project these four patterns. Open-model time series (OMTS) techniques analyze the time series to determine which patterns exist and then build a unique model of that time series to project the patterns into the future and, thus, to forecast the time series. This is in contrast to fixed-model time series (FMTS) techniques, which have fixed equations that are based upon *a priori* assumptions that certain patterns do or do not exist in the data.

In fact, when you consider both OMTS and FMTS techniques, there are more than 60 different techniques that fall into the general category of time series techniques. Fortunately, we do not have to explain each of them in this chapter. This is because some of the techniques are very sophisticated and take a considerable amount of data but do not produce any better results than simpler techniques, and they are seldom used in practical sales forecasting situations. In other cases, several different time series techniques may use the same approach to forecasting and have the same level of effectiveness. In these latter cases where several techniques work equally well, we will discuss only the one that is easiest to understand (following the philosophy, why make something complicated if it does not have to be). This greatly reduces the number of techniques that need to be discussed.

Because they are generally easier to understand and use, we will start with FMTS techniques and return to OMTS later in the chapter.

❖ FIXED-MODEL TIME SERIES TECHNIQUES

FMTS techniques are often simple and inexpensive to use and require little data storage. Many of the techniques (because they require little data) also adjust very quickly to changes in sales conditions and, thus, are appropriate for short-term forecasting. We can fully understand the range of FMTS techniques by starting with the concept of an average as a forecast (which is the basis on which all FMTS techniques are founded) and move through the levels of moving average, exponential smoothing, adaptive smoothing, and incorporating trend and seasonality.

The Average as a Forecast

All FMTS techniques are essentially a form of average. The simplest form of an average as a forecast can be represented by the following formula:

$$\text{Forecast}_{t+1} = \text{Average Sales}_{1 \text{ to } t} = \sum_{t=1}^{N} S_t / N \qquad (1)$$

where: S = Sales
N = Number of Periods of Sales Data (t)

In other words, our forecast for next month (or any month in the future, for that matter) is the average of all sales that have occurred in the past.

The advantage to the average as a forecast is that the average is designed to "dampen" out any fluctuations. Thus, the average takes the noise (which time series techniques assume cannot be forecast anyway) out of the forecast. However, the average also dampens out of the forecast **any** fluctuations, including such important fluctuations as trend and seasonality. This principle can be demonstrated with a couple of examples.

Figure 3.2 provides a history of sales that has only the time series components of level and noise. The forecast (an average) does a fairly good job of ignoring the noise and forecasting only the level. However, Figure 3.3 illustrates a history of sales that has the time series components of level and noise, *plus trend.* As will always happen when

Figure 3.2 Average as a Forecast: Level and Noise

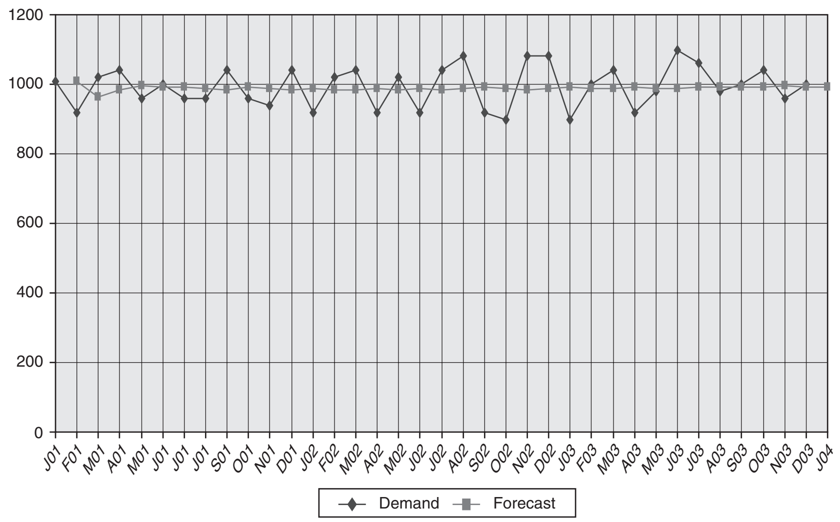

the average is used to forecast data with a trend, the forecast always lags behind the actual data. Because the average becomes more "sluggish" as more data are added, the lagging of the forecast behind the actual sales gets worse over time. If our example in Figure 3.3 had been a negative trend, lagging behind would have meant the average would have always forecast high.

As a final example, Figure 3.4 illustrates a history of sales that has the time series components of level and noise, *plus seasonality.* Notice that the average has the unfortunate effect of losing (dampening out) the seasonal pattern. Thus, we would lose this important component of any possible forecast.

The conclusion from these three illustrations is that the average should only be used to forecast sales patterns that contain only the time series components of level and noise. Remember that FMTS techniques assume certain patterns exist in the data. In the case of the average, we are assuming there is no trend or seasonality in the data. This is why we stated earlier that the forecast for the next period is also the forecast for all future periods. Because the data are supposed to be level, there should be no pattern of sales increasing (trend) or increasing and

Figure 3.3 Average as a Forecast: Level, Trend, and Noise

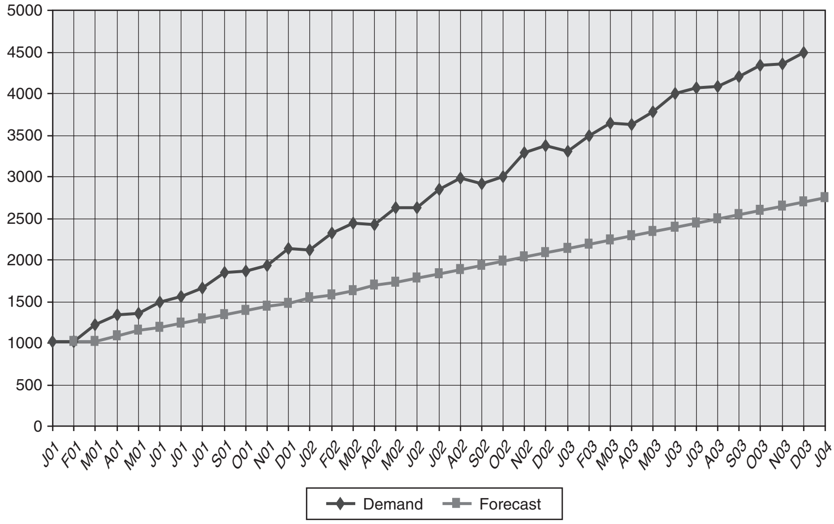

Figure 3.4 Average as a Forecast: Level, Seasonality, and Noise

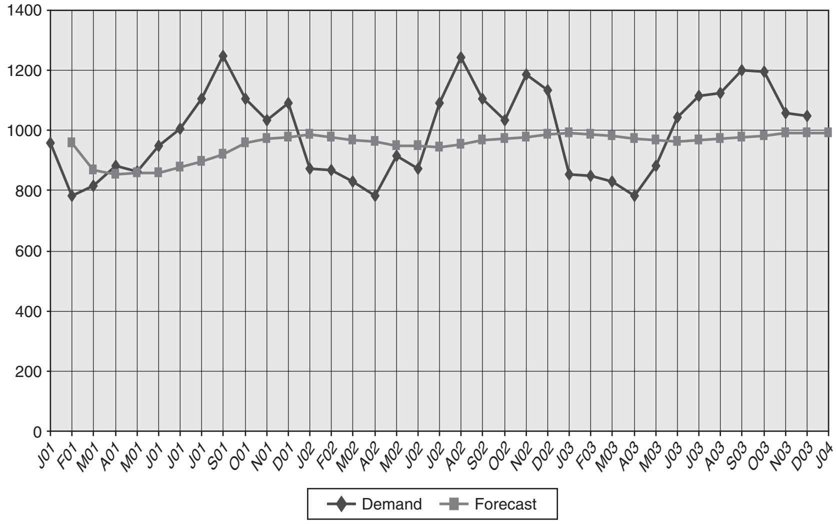

decreasing (seasonality). Therefore, sales should be the same (level) for each period in the future. If nothing else, this demonstrates the rather naïve assumption that accompanies the use of the average as a forecast.

The average as a forecasting technique has the added disadvantage that it requires an ever-increasing amount of data storage. With each successive month, an additional piece of data must be stored for the calculation. With the data storage capabilities of today's computers, this may not be too onerous a disadvantage, but it does cause the average to be sluggish to changes in level of demand. One last example should illustrate this point. Figure 3.5 shows a data series with little noise, but the level changes. Notice that the average as a forecast never really adjusts to this new level because we cannot get rid of the "old" data (the data from the previous level).

Thus, the average as a forecast does not consider trend or seasonality, and it is sluggish to react to changes in the level of sales. In fact, it does little for us as a forecasting technique, other than give us an excellent starting point. All FMTS techniques were developed to overcome some disadvantage of the average as a forecast. We next explore the first attempt at improvement, a moving average.

Figure 3.5 Average as a Forecast: Level Change

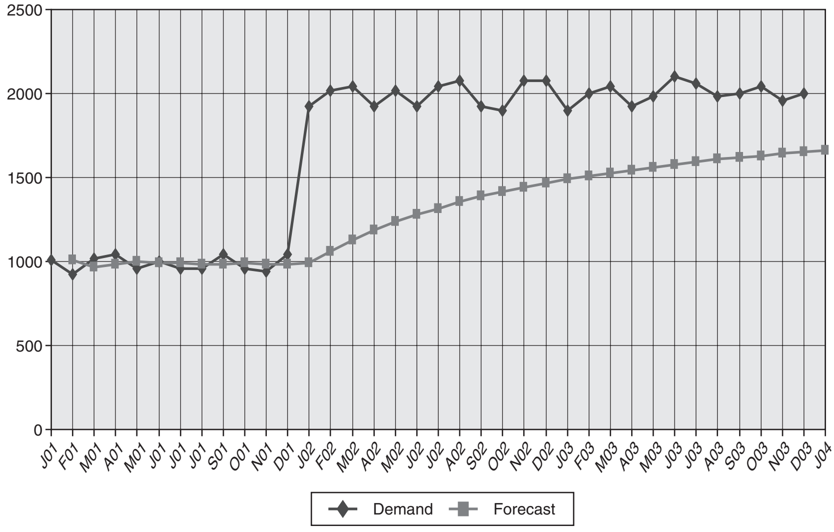

Moving Average

Rather than use all the previous data in the calculation of an average as the forecast, why not just use some of the more recent data? This is precisely what a moving average does, with the following formula.

$$F_{t+1} = (S_t + S_{t-1} + S_{t-2} + \cdots + S_{t-N-1})/N \tag{2}$$

where: F_{t+1} = Forecast for Period t + 1
S_{t-1} = Sales for Period t − 1
N = Number of Periods in the Moving Average

So a three-period moving average would be:

$$F_{t+1} = (S_t + S_{t-1} + S_{t-2})/3$$

a four-period moving average would be:

$$F_{t+1} = (S_t + S_{t-1} + S_{t-2} + S_{t-3})/4$$

a five-period moving average would be:

$$F_{t+1} = (S_t + S_{t-1} + S_{t-2} + S_{t-3} + S_{t-4})/5$$

and so on, for as many periods in the moving average as you would like.

The problem with a moving average is deciding how many periods of sales to use in the forecast. The more periods used, the more it starts to look like an average. The fewer periods used, the more reactive the forecast becomes, but the more it starts to look like our naïve technique from Chapter 2 (the forecast for the next period equals the sales from the last period). Applying 3-period, 6-period, and 12-period moving averages to each of the demand patterns in Figures 3.2, through 3.5 (now Figures 3.6 through 3.9, respectively) should illustrate some of these points.

For a time series that has only level and noise (Figure 3.6), our three moving averages work equally well. This is because all dampen out the relatively small amount of noise, and there is no change in level to which to react. Because it uses the least data, the three-period moving average is superior in this case.

However, for the time series with trend added (Figure 3.7), very different results are obtained. The longer the moving average, the less reactive the forecast, and the more the forecast lags behind the trend (because it is more like the average). Again, this is because moving averages were not really designed to deal with a trend, but the shorter moving averages adjust better (are more reactive) than the longer in this case.

An interesting phenomenon occurs when we look at the use of moving averages to forecast time series with seasonality (Figure 3.8). Notice that both the three-period and the six-period moving averages lag behind the seasonal pattern (forecast low when sales are rising and forecast high when sales are falling) and miss the turning points in the time series. Notice also that the more reactive moving average (three-period) does a better job of both of these. This is because in the short run (defined here as between turning points), the seasonal pattern simply looks like trend to a moving average.

However, the 12-period moving average simply ignores the seasonal pattern. This is due to the fact that any average dampens out

Figure 3.6 Moving Average as a Forecast: Level and Noise

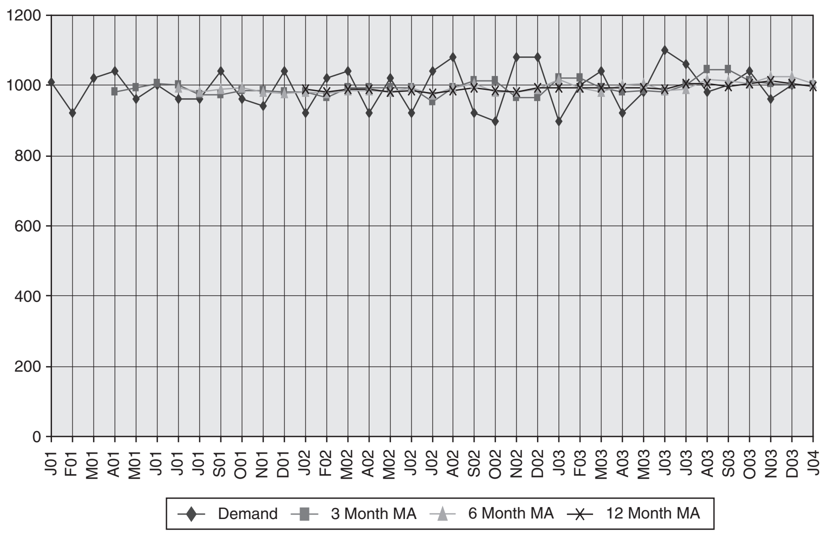

Figure 3.7 Moving Average as a Forecast: Level, Trend, and Noise

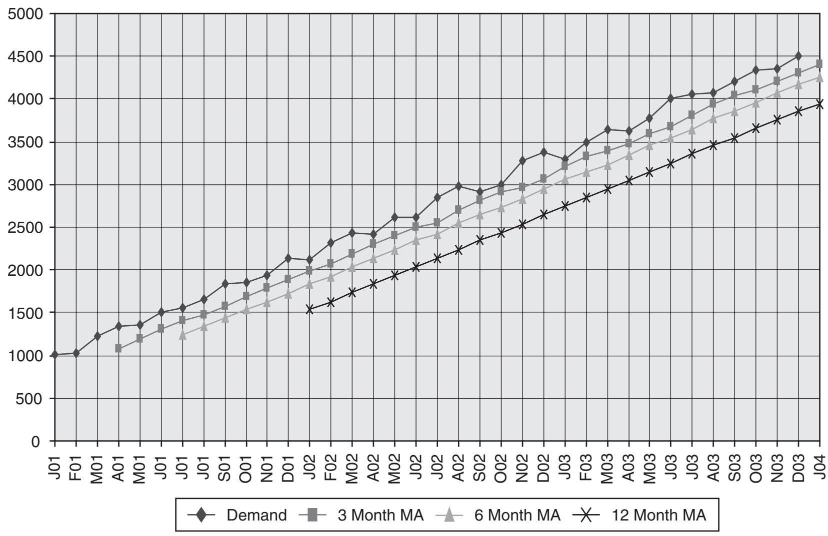

Figure 3.8 Moving Average as a Forecast: Level, Seasonality, and Noise

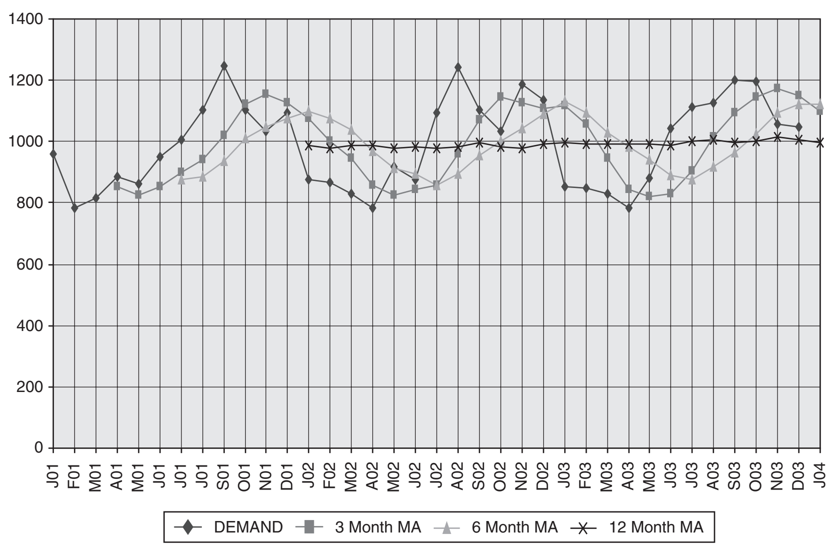

Figure 3.9 Moving Average as a Forecast: Level Changes

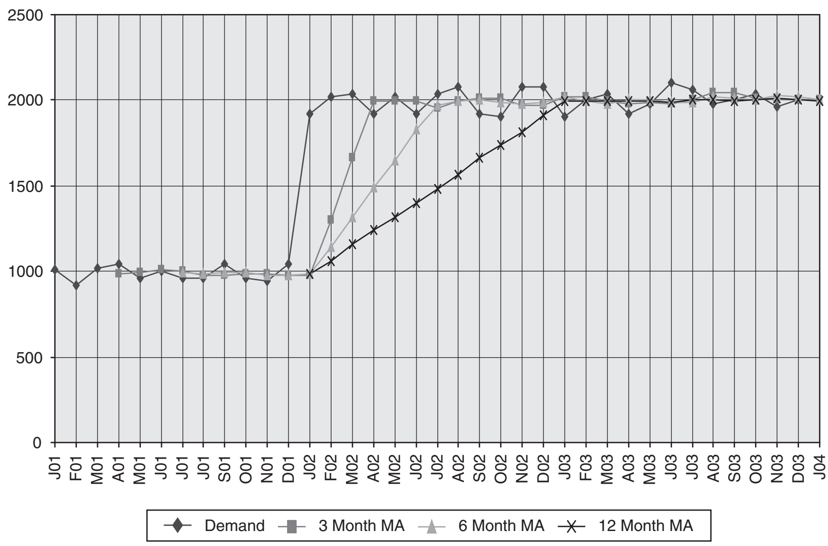

random fluctuations (noise) ***and*** any patterns that are the same length as the average. Because this time series has a 12-month seasonal pattern, a 12-month moving average completely loses the seasonal component in its forecast. This is particularly dangerous when you consider how many sales managers use a simple 12-month moving average to generate a forecast—they are inadvertently dampening out the seasonal fluctuations from their forecasts.

Finally, let's look at the time series where the level changes (Figure 3.9). Again, the longer moving average tends to dampen out the noise better than the shorter moving average, but the shorter moving average reacts more quickly to the change in level.

Thus, what we need in a moving average is one that acts like an average when there is only noise in the time series (dampens out the noise but uses less data than an average), but acts like a naïve forecast when the level changes (puts more weight on what happened very recently). The problem with this is how to recognize the difference in a change that is noise, as opposed to a change in level, a trend, or a seasonal pattern.

A final problem with the moving average is that the same weight is put on all past periods of data in determining the forecast. It is more reasonable to put greater weight on the more recent periods than the older periods (especially when a longer moving average is used). Therefore, the question when using a moving average becomes how many periods of data to use and how much weight to put on each of those periods. To answer this question about moving averages, a technique called exponential smoothing was developed.

Exponential Smoothing

Exponential smoothing is the basis for almost all FMTS techniques in use today. It is easier to understand this technique if we acknowledge that it was originally called an "exponentially weighted moving average." Obviously, the original name was too much of a mouthful for everyday use, but it helps us to explain how this deceptively complex technique works. We are going to develop a moving average, but we will weight the more recent periods of sales more heavily in the forecast, and the weights for the older periods will decrease at an exponential rate (which is where the "exponential smoothing" term came from).

Regardless of that rather scary statement, we are going to accomplish this with a very simple calculation (Brown & Meyer, 1961).

$$F_{t+1} = \alpha S_t + (1 - \alpha) F_t \tag{3}$$

where: F_t = Forecast for Period t
S_t = Sales for Period t
$0 < \alpha < 1$

In other words, our forecast for next period (or, again, any period in the future) is a function of last period's sales and last period's forecast, with this α thing thrown in to confuse us.

What we are actually doing with this exponential smoothing formula is merely a weighted average. Because α is a positive fraction (that is, between 0 and 1), $1 - \alpha$ is also a positive fraction, and the two of them add up to 1. Any time we take one number and multiply it by a positive fraction, take a second number and multiply it by the reciprocal of the positive fraction (another way of saying 1 – the first fraction), and add the two results together, we have merely performed a weighted average. Several examples should help:

1. When we want to average two periods' sales (Period 1 was 50 and Period 2 was 100, for example) and not put more weight on one than the other, we are actually calculating it as $((0.5 \times 50) + (0.5 \times 100)) = 75$. We simply placed the same weight on each period. Notice that this gives us the same result as if we had done the simpler equal-weight average calculation of $(50 + 100)/2$.

2. When we want the same two periods of sales but want to put three times as much weight on Period 2 (for reasons we will explain later), the calculation would now be $((0.25 \times 50) + (0.75 \times 100)) = 87.5$. Notice that in this case α would be 0.25 and $1 - \alpha$ would be 0.75.

3. Finally, if we want nine times as much weight on Period 2, the resultant calculation would be $((0.1 \times 50) + (0.9 \times 100)) = 95$. Again, notice that in this case α would be 0.1 and $1 - \alpha$ would be 0.9.

Therefore, we can control how much emphasis in our forecast is placed on what sales actually were last period. But what is the purpose of using last period's forecast as part of next period's forecast? This is where exponential smoothing is "deceptively complex" and requires some illustration.

For the purpose of this illustration, let's assume that on the evening of the last day of each month, we make a forecast for the next month. Let's also assume that we have decided to use exponential smoothing and to put 10% of the weight of our forecast on what happened last month. Further, let's assume this is the evening of the last day of June. Thus, our value for α would be 0.1 and our forecast for July would be:

$$F_{JULY} = .1\ S_{JUNE} + .9\ F_{JUNE}$$

But where did we get the forecast for June? In fact, a month ago on the evening of the last day of May, we made this forecast:

$$F_{JUNE} = .1\ S_{MAY} + .9\ F_{MAY}$$

Again, where did we get the forecast for May? And again, a month ago on the evening of the last day of April, we made this forecast:

$$F_{MAY} = .1\ S_{APRIL} + .9\ F_{APRIL}$$

We could keep this up forever, but suffice it to say that each month the forecast from the previous month has in it the forecasts (and the sales) from all previous months. Thus, 10% of the forecast for July is made up of sales from June, but the other 90% is made up of the forecast for June. However, the forecast for June was made up of 10% of the sales from May. Thus, 90% times 10% (or 9%) of the July forecast is made up of the sales from May. The rest of the forecast for June was made up of 90% of the forecast for May, which in turn was made up of 10% of the sales from April (so April sales comprises 90% times 90% times 10%, or 8.1%, of the July forecast) and 90% of the forecast from April, and so on back to where we made our first forecast. This leads us to the fact that the forecast for July is actually made up of the following rather complicated formula:

$$\begin{aligned} F_{JULY} = .1\ S_{JUNE} + (.9)\ (.1)\ S_{MAY} + (.9)^2\ (.1)\ S_{APRIL} \\ + (.9)^3\ (.1)\ S_{MARCH} + \cdots + (.9)^N\ (.1)\ S_{JULY-(N+1)} \end{aligned}$$

If we take a second to study this formula, we see that sales from June make up 10% of our forecast, sales from May make up 9% (.9 × .1) of our forecast, sales from April make up 8.1% (.9 × .9 × .1) of our forecast, sales from March make up 7.2% (.9 × .9 × .9 × .1) of our forecast, and so on back to the first month we used this technique.

What is happening with the rather simple-looking exponential smoothing formula is that we are putting α weight on last period's sales, α times $(1 - \alpha)$ weight on the previous period's sales, and changing the weight for each previous period's sales by multiplying the weight by $(1 - \alpha)$ for each successive period we go into the past.

For $\alpha = 0.1$, this causes the weights for the previous period's sales to decrease at the following exponential rate: 0.1, 0.09, 0.081, 0.072, 0.063, . . . and for $\alpha = 0.2$, the weights for the previous period's sales to decrease at the following exponential rate: 0.2, 0.16, 0.128, 0.1024, 0.08192, . . .

We could try to develop a similar series for every value of α (by the way, the possible values of α between 0 and 1 are infinite, so our attempt might take a while), but it is not necessary—the simple exponential smoothing formula does it for us. We do need to remember, however, that the higher the value of α, the more weight we are putting on last period's sales and the less weight we are putting on all the previous periods combined. In fact, as α approaches one, exponential smoothing puts so much weight on the past period's sales and so little on the previous periods combined, that it starts to look like our naïve technique ($F_{t+1} = S_t$) from Chapter 2. Conversely, as α approaches zero, exponential smoothing puts more equal weight on all periods and starts to look much like the average as a forecast.

This leads us to some conclusions about what the value of α should be:

1. The more the level changes, the larger α should be, so that exponential smoothing can quickly adjust.
2. The more random the data, the smaller α should be, so that exponential smoothing can dampen out the noise.

Several examples should help illustrate these conclusions. For our first illustration, we can use the data pattern from Figure 3.9 for the moving average, now Figure 3.10 for exponential smoothing.

Figure 3.10 Exponential Smoothing as a Forecast: Level Change

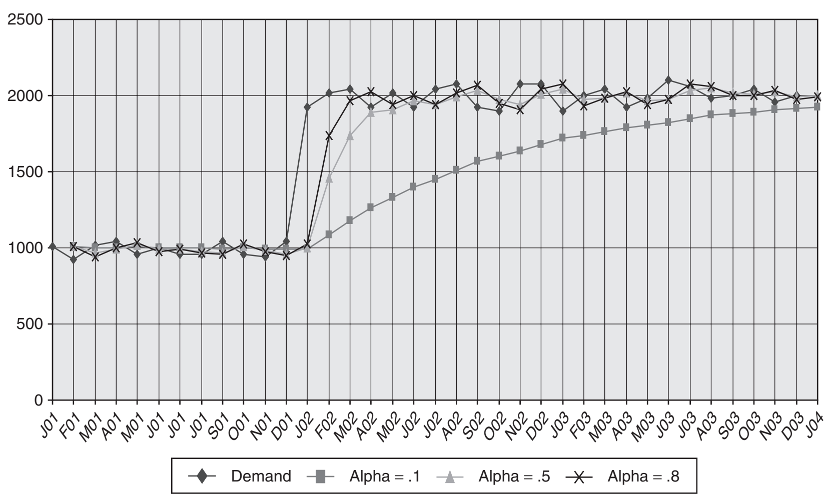

In Figure 3.10, we can see three exponential smoothing forecasts of the time series. All three do a fairly good job when the level is stable, but the higher the value of α in the forecast, the quicker it reacts to the change in level. Because a low value of α is much like an average, the forecast for the low α never quite reaches the new level.

However, a very different result is found when we observe the forecasts of the time series in Figures 3.11 and 3.12. Figure 3.11 is a reproduction of the data series used in Figures 3.2 and 3.6 and represents a time series with no trend and a low amount of noise. In this series, the exponential smoothing forecasts with various levels of α all perform fairly well. However, in the time series of Figure 3.12, which has a stable level but a high amount of noise, the forecasts with the higher values of α overreact to the noise and, as a result, jump around quite a bit. The forecast with the lower level of α does a better job of dampening out the noise.

Given these illustrations of our conclusions about the value of α that should be used, we have in exponential smoothing a technique that overcomes many of the problems with the average and the moving average as forecasting techniques. Exponential smoothing is less

Figure 3.11 Exponential Smoothing as a Forecast: Low Noise

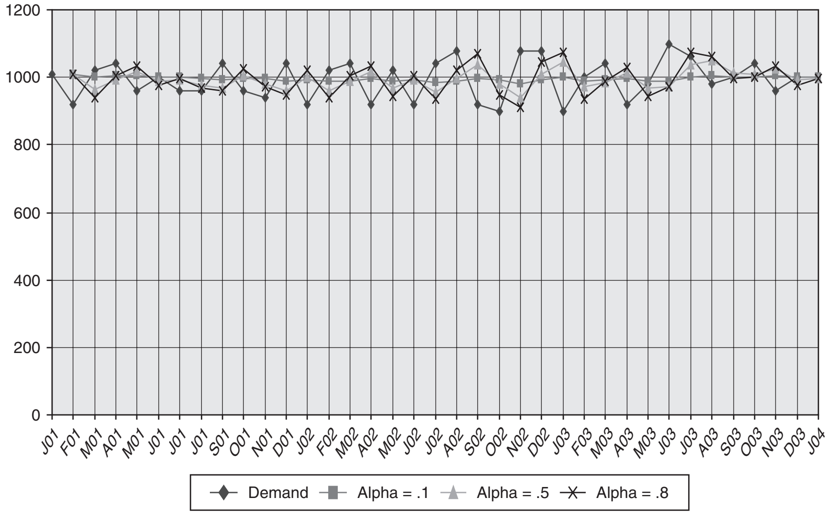

Figure 3.12 Exponential Smoothing Average as a Forecast: High Noise

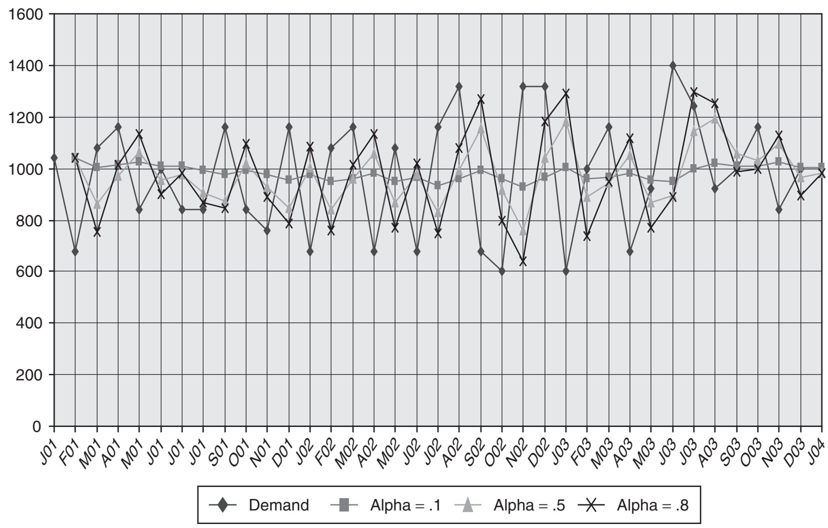

cumbersome than the average because exponential smoothing only requires the values of last period's sales and forecast and the value of α. Exponential smoothing solves the problems with the moving average of how much data to use and how to weight it by using an exponentially decreasing weight for all previous periods.

However, with exponential smoothing we are still faced with a dilemma: How do we determine whether the level is changing or if it is simply noise and, thus, what the value of α should be? To answer this dilemma, the next group of techniques (called adaptive smoothing) was developed.

Adaptive Smoothing

Although a number of adaptive smoothing techniques exist, they all have one thing in common: each is an attempt to automatically select the value of α. Because there are so many adaptive smoothing techniques and they all work essentially equally well, we will only discuss the simplest of this group of techniques here. This adaptive smoothing approach uses the absolute value of the percent error from the previous period's forecast to adjust the value of α for the next period's forecast (Trigg & Leach, 1967). Thus, the original exponential smoothing formula is still used:

$$F_{t+1} = \alpha\, S_t + (1 - \alpha)\, F_t \tag{4}$$

but after each period's sales are recorded, the value of α is adjusted for the next period by the following formula:

$$\alpha_{t+2} = \mid (F_{t+1} - S_{t+1})/S_{t+1} \mid = \mid PE_{t+1} \mid \tag{5}$$

Because Equation (5) can produce values outside the range of α, this calculation is adjusted by the following rules:

If $\mid PE_{t+1} \mid$ is equal to or greater than 1.0, then $\alpha_{t+2} = 0.99999$

If $\mid PE_{t+1} \mid$ is equal to 0.0, then $\alpha_{t+2} = 0.00001$

We can illustrate the adaptability of this technique by forecasting the times series with level change in Figure 3.10, now Figure 3.13 for

Figure 3.13 Adaptive Smoothing as a Forecast: Level Change

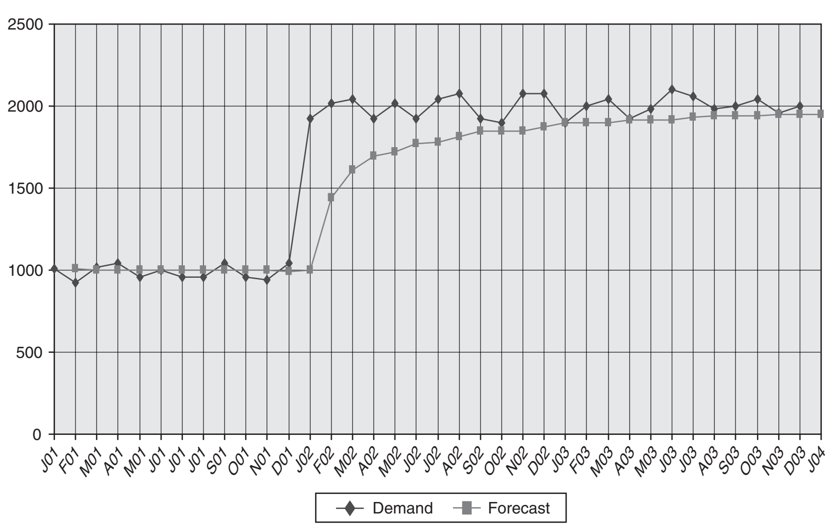

adaptive smoothing. To illustrate the changes in α that result in this technique, the calculations are also reproduced in Table 3.1.

To get the process started, we used the usual convention of setting the initial value of α at 0.1, although any value can be chosen without changing the resultant forecasts. The reason for this is that we also assume that the initial forecast was equal to the first period demand, so the first forecast becomes:

$$F_2 = \alpha S_1 + (1 - \alpha) S_1$$

So regardless of the initial value of α that is chosen, the forecast for period two is always equal to sales from period one. The true calculation of a forecast and the adapted values of α begin at that point.

Notice that the value of α stays low (well below 0.1) while the time series is level (a low value of α dampens out the noise), but as soon as the level changes, the value of α jumps dramatically to adjust. Once the time series levels off, the value of α again returns to a low level.

This adaptive smoothing technique overcomes one of the major problems with exponential smoothing: what should be the value chosen for α? However, all the techniques we have discussed so far have a common problem: none of them considers trend or seasonality. Since this technique assumes there is no trend or seasonality, our forecast of January 2004 is 1950 and is also our forecast for *every month* in 2004—we assume there will be no general increase or decrease in sales (trend), nor will there be any pattern of fluctuation in sales (seasonality). Because this is unrealistic for many business demand situations, we need some way to incorporate trend and seasonality into our FMTS forecasts. To do so, we temporarily set aside the concept of smoothing constant adaptability and introduce first trend and then seasonality into our exponential smoothing calculations.

Exponential Smoothing With Trend

Although we tend to think of trend as a straight or curving line going up or down, for the purposes of exponential smoothing, it is helpful to think of trend as a series of changes in the level. In other words, with each successive period, the level either "steps up" or "steps down." This "step function," or changing level pattern, of trend is conceptually illustrated in Figure 3.14. Although demand is going up

Table 3.1 Adaptive Smoothing Forecast Calculations

Month	*Demand*	*Forecast*	*Percent Error*	*Absolute PE or* α_{t+1}
J01	1010			
F01	920	1010	0.098	0.098
M01	1020	1002	−0.018	0.018
A01	1040	1002	−0.037	0.037
M01	960	1003	0.045	0.045
J01	1000	1001	0.001	0.001
J01	960	1001	0.043	0.043
A01	960	999	0.041	0.041
S01	1040	998	−0.041	0.041
O01	960	999	0.041	0.041
N01	940	998	0.061	0.061
D01	1040	994	−0.044	0.044
J02	1920	996	−0.481	0.481
F02	2020	1441	−0.287	0.287
M02	1920	1607	−0.163	0.163
A02	2040	1658	−0.187	0.187
M02	2080	1729	−0.169	0.169
J02	1920	1789	−0.068	0.068
J02	2040	1798	−0.119	0.119
A02	2080	1826	−0.122	0.122
S02	1920	1857	−0.033	0.033
O02	1900	1859	−0.021	0.021
N02	2080	1860	−0.106	0.106
D02	2080	1883	−0.095	0.095
J03	1900	1902	−0.001	0.001
F03	2000	1902	−0.049	0.049
M03	2040	1907	−0.065	0.065
A03	1920	1916	−0.002	0.002
M03	1980	1916	−0.033	0.033
J03	2100	1918	−0.087	0.087
J03	2060	1933	−0.061	0.061
A03	1980	1941	−0.020	0.020
S03	2000	1942	−0.029	0.029
O03	2040	1944	−0.047	0.047
N03	1960	1948	−0.006	0.006
D03	2000	1948	−0.026	0.026
J04		1950		

Figure 3.14 Trend in a Time Series

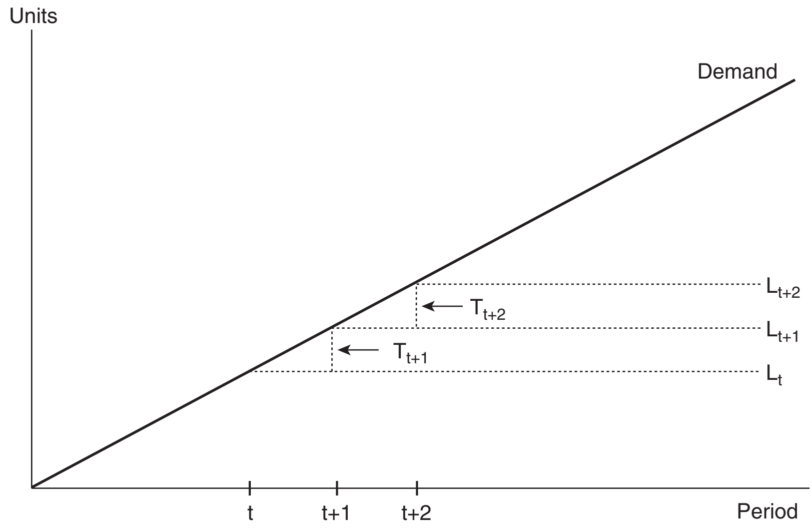

in a straight line, we can conceive of it as a series of increases in the level (the dashed horizontal lines). This is much like climbing a set of stairs. Although we make steady progress up the stairs, we are actually stepping up one step each period (the amount we step up, or the height of each step, on a set of stairs is called the riser). The height of each step (the riser) is what we call "trend" in exponential smoothing, and that trend is designated in Figure 3.14 as T. For period t + 1, the trend is the amount the level changed from period t to period t + 1 ($L_{t+1} - L_t$), or T_{t+1}. Similarly for period t + 2, the trend is the amount the level changed from period t + 1 to period t + 2 ($L_{t+2} - L_{t-1}$), or T_{t+2}.

To understand the calculation of trend in exponential smoothing, we must also understand that an exponential smoothing calculation is just a weighted average of two measures of the same thing. Our original exponential smoothing formula (Equation 4) was:

$$F_{t+1} = \alpha\, S_t + (1 - \alpha)\, F_t$$

In this calculation, S_t is one measure of past sales (last period's sales) and F_t is another measure of past sales (a weighted average of sales in all periods prior to t). Thus, we were taking a weighted average of two measures of the same thing. We are now going to do the same thing for level and trend with the following formulae (Holt et al., 1960):

$$L_t = \alpha\, S_t + (1 - \alpha)\, (L_{t-1} + T_{t-1}) \tag{6}$$

$$T_t = \beta\, (L_t - L_{t-1}) + (1 - \beta)\, T_{t-1} \tag{7}$$

where: L = Level
T = Trend
$0 < \alpha < 1$
$0 < \beta < 1$

Notice that Equation (6) looks very similar to our earlier exponential smoothing forecast calculation—we still use α in the same way and we still use last period's sales. This difference is the addition of trend into the second part of Equation (6) and the fact that it is not a forecast for next period (F_{t+1}) but rather a measure of level for this period (L_t). In fact, in our original exponential smoothing formula (Equation [4]), we did not include trend because we assumed it did not exist. Because trend was assumed not to exist, our estimate of level this period ***was*** our forecast of next period.

What we need are two estimates of level for this period so we can exponentially smooth them. The first estimate is simply sales for this period. Since we assume there is no seasonality in the time series (an assumption we will discard in the next section), then this sales value has no seasonality in it. Because the trend is a change in level from one period to the next, any given value of sales does not have trend in it (that is, trend is in the *change* in sales from one period to the next, not any single sales value). Finally, when we perform the weighted averaging of the exponential smoothing calculation in Equation (6), we get rid of the noise. (Remember that averaging removes noise.) Since this logic says there is no trend or seasonality in the sales value and we will get rid of this noise when we do our exponential smoothing calculation, we are only left with one time series component in the sales value, and that component is level.

The second estimate of level is our estimate of level from last period, plus the estimate of how much level should have changed from last period to this period (that is, the trend). This gives us two measures of level to exponentially smooth with α.

Our two estimates of trend in Equation (7) are how much the level changed from last period to this period and our estimate of trend from last period. These two measures of trend are exponentially smoothed with our new smoothing constant, β. β is just like α in that it is a positive fraction (that is, between zero and one). It is designated by a different Greek symbol to indicate that α and β can have different values.

Once we have our new estimates of level (L) and trend (T), we can forecast as far into the future as we want by taking the level and adding to it the trend per period times as many periods into the future as we want the forecast. This can be represented by the following formula:

$$F_{t+m} = L_t + (T_t \times m) \tag{8}$$

where: m = the number of periods into the future to forecast.

To illustrate this technique, consider the time series with trend introduced in Figure 3.3, now Figure 3.15 for exponential smoothing with trend. To illustrate the calculations involved in this technique, Table 3.2 provides the calculations of level and trend and a forecast forward for one period throughout the time series (also provided in Figure 3.15). For the purposes of illustration, we arbitrarily chose the value of 0.1 for α and the value of 0.2 for β. Notice that to get the process started, we used the usual convention of assuming the level for the first period equaled first period demand, and the trend for the first period equaled the change in demand from the first to the second period.

To provide a forecast for any period more than one in the future (April 2004, for example), it is merely a task of taking the most recent value of level that has been calculated (in this case, December 2003) and adding to it the most recent value of trend that has been calculated (also in this case, December 2003) times the number of months into the future that we wish to forecast (because April is four months past December, it would be times four). For April 2004, the calculations are:

$$F_{A04} = L_{D03} + (T_{D03} \times m)$$

$$F_{A04} = 4614 + (106 \times 4)$$

$$F_{A04} = 5038$$

Figure 3.15 Exponential Smoothing With Trend as a Forecast

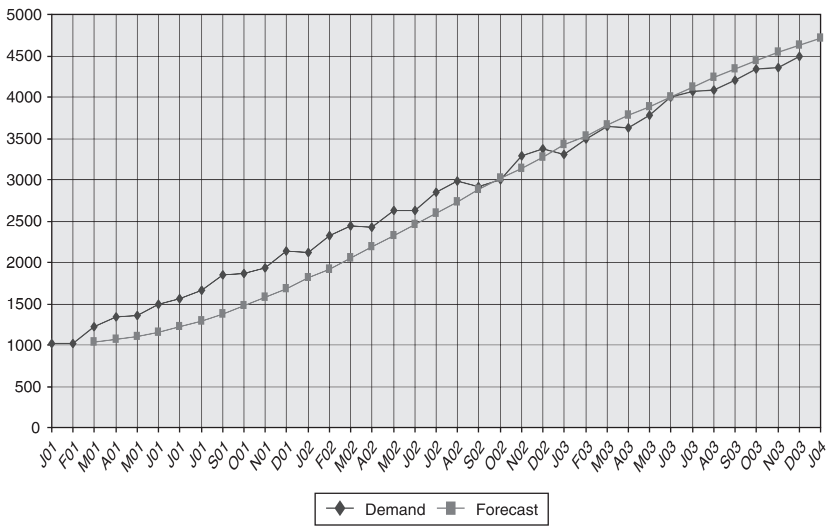

Now that we have the logic for introducing trend into the exponential smoothing calculations, it is fairly easy to also bring in seasonality.

Exponential Smoothing With Trend and Seasonality

To introduce seasonality, let's first think of a simple demand example where we sell 12,000 units of a product every year. If there is no trend, no noise, and no seasonality, we would expect to sell 1,000 units every month (that is, the level). If, however, we noticed that every January we sold, on average, 1,150 units, there is clearly a pattern here of selling more than the level in January. In fact, we are selling 1,150/1,000, or 1.15, times the level.

This value of 1.15 is called a *multiplicative seasonal adjustment* and means that sales in that month are 15% higher than they would be without a seasonal pattern. Similarly, a seasonal adjustment of 1.00 means that sales are right at the non-seasonal level, and a seasonal adjustment of 0.87 means that sales are 13% below what we would expect if there was no seasonal pattern.

Table 3.2 Exponential Smoothing With Trend Forecast Calculations

Month	*Demand*	*Level* *($\alpha = 0.1$)*	*Trend* *($\beta = 0.2$)*	*Forecast*
J01	1010	1010	10	
F01	1020	1020	10	
M01	1220	1049	14	1030
A01	1340	1091	19	1063
M01	1360	1135	24	1110
J01	1500	1193	31	1159
J01	1560	1258	38	1224
A01	1660	1332	45	1296
S01	1840	1424	54	1377
O01	1860	1516	62	1478
N01	1940	1615	69	1578
D01	2140	1729	78	1684
J02	2120	1839	85	1808
F02	2320	1963	93	1924
M02	2440	2094	100	2056
A02	2420	2217	105	2195
M02	2620	2352	111	2322
J02	2620	2478	114	2462
J02	2840	2617	119	2592
A02	2980	2760	124	2736
S02	2920	2887	124	2884
O02	3000	3011	124	3012
N02	3280	3149	127	3135
D02	3380	3287	129	3277
J03	3300	3404	127	3416
F03	3500	3528	126	3531
M03	3640	3653	126	3654
A03	3620	3763	123	3779
M03	3780	3875	121	3886
J03	4000	3996	121	3996
J03	4060	4111	120	4117
A03	4080	4216	117	4231
S03	4200	4319	114	4332
O03	4340	4424	112	4433
N03	4360	4518	109	4536
D03	4500	4614	106	4627
J04				4720

We are now going to use this concept of a multiplicative seasonal adjustment to introduce seasonality into the exponential smoothing calculations. Again, we will develop two different measures of each seasonal adjustment and take a weighted average of them (through exponential smoothing) to come up with our new estimate. To do this, however, we also need to update our formulae for Exponential Smoothing with Trend (Equations [6] and [7]) to take into account the fact that seasonality is now assumed to exist. This leads us to the following formulae (Winters, 1960):

$$L_t = \alpha\,(S_t/SA_{t-C}) + (1-\alpha)\,(L_{t-1} + T_{t-1}) \tag{9}$$

$$T_t = \beta\,(L_t - L_{t-1}) + (1-\beta)\,T_{t-1} \tag{10}$$

$$SA_t = \gamma\,(S_t/L_t) + (1-\gamma)\,(SA_{t-C}) \tag{11}$$

where: L = Level
T = Trend
SA_t = Seasonal Adjustment for Period t
C = The Cycle Length of the Seasonal Pattern (that is, the cycle length for a 12-month pattern is C = 12)
$0 < \alpha < 1$
$0 < \beta < 1$
$0 < \gamma < 1$

We have revised our calculation for level to take seasonality into account in our first estimate of level. Recalling our previous example of annual sales of 12,000, how would we take the seasonality out of January sales? If sales were 1,150 and our previous estimates of the seasonality adjustment for January were 1.15, we can de-seasonalize January sales simply by dividing the sales value of 1,150 by the seasonal adjustment of 1.15. This gives us a de-seasonalized value of 1,000—precisely the value we said was the expected level if there were no seasonality.

Thus, by dividing sales for any period by the seasonal adjustment for the same period last year (that is, divide sales for January 2004 by the seasonal adjustment for January 2003), we have an estimate in the first formula of level with the seasonality taken out (recall that the original formula already took out the trend and the noise). Because the second part of this formula contains the level from the last period, which was de-seasonalized at that time, we now have two estimates of level to exponentially smooth.

Thankfully, the formula for trend (Equation [10]) does not change. Therefore, we do not have to revisit it here.

However, we now have added a formula to calculate the seasonal adjustments (Equation [11]). Again, we need two estimates of the seasonal adjustment for each period, so we can exponentially smooth each. This means we have 12 of these calculations per year if we are forecasting monthly sales, 52 if we are forecasting weekly, and 4 if we are forecasting quarterly.

The first part of Equation (11) is, again, a throwback to our initial example. If we take the sales value for this period and divide it by the most recently calculated level (which was just done two formulas before and is L_t), we have one estimate of the seasonal adjustment for this period. In our initial example, we did the same thing when we divided 1,150 by 1,000 to obtain 1.15 as our estimate of the seasonal adjustment for January.

For our second estimate of the seasonal adjustment for this period, we need to look back one year to the same period last year. We can now exponentially smooth these two estimates of the seasonal adjustment for this period using the smoothing constant, γ. Again, γ is just like α and β in that it is a positive fraction (that is, between zero and one). It is designated by a different Greek symbol to indicate that α, β, and γ can all have different values.

Once we have our new estimates of level (L), trend (T), and seasonal adjustments (SA), we can forecast as far into the future as we want by taking the level, adding to it the trend per period times as many periods into the future as we want the forecast, and multiplying that result by the most recent seasonal adjustment for that period. This can be represented by the following formula:

$$F_{t+m} = (L_t + (T_t \times m)) \times SA_{t-C+m} \tag{12}$$

where: m = the number of periods into the future to forecast.

The last component of Equation (12) probably needs a little illustration. If we have just received sales for December 2003 and want to forecast April 2004, we will use the values of L and T calculated in December 2003 (in this case, December 2003 is "t") for the first part of the forecast. However, our most recent estimate of the seasonal adjustment for April was calculated back in

April 2003. The symbol to represent using this value is to take "t" (December 2003); subtract C, or 12, months from it (placing us in December 2002); and add to it m, or 4, months to bring us to the seasonal adjustment for April 2003.

To illustrate this technique, we will go all the way back to our original time series with trend and seasonality introduced in Figure 3.1, now Figure 3.16, for exponential smoothing with trend and seasonality. To illustrate the calculations involved in this technique, Table 3.3 provides the calculation of level, trend, seasonality, and a forecast forward for one period throughout the time series (also provided in Figure 3.16). For the purposes of illustration, we arbitrarily chose the value of 0.1 for α, the value of 0.2 for β, and the value of 0.15 for γ. Notice that to get the process started, we used the usual convention of assuming the level for the first period equaled first period demand, the trend for the first period equaled the change in demand from the first to the second period, and the initial 12 seasonal adjustment values were equal to 1.00. Notice, also, that this technique does a pretty terrible job

Figure 3.16 Exponential Smoothing With Trend and Seasonality as a Forecast

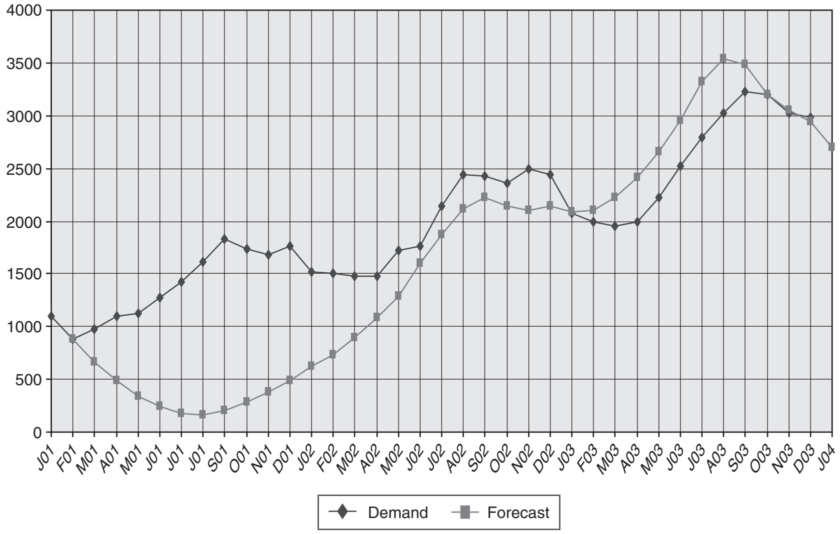

Table 3.3 Exponential Smoothing With Trend and Seasonality Forecast Calculations

Month	*Demand*	*Level ($\alpha = 0.1$)*	*Trend ($\beta = 0.2$)*	*Seasonality ($\gamma = 0.15$)*	*Forecast*
J01	1104	1104	−219	1.00	
F01	885	885	−219	1.00	885
M01	976	697	−213	1.06	666
A01	1101	546	−200	1.15	484
M01	1120	423	−185	1.25	345
J01	1276	342	−164	1.41	238
J01	1419	302	−139	1.56	178
A01	1615	308	−110	1.64	162
S01	1836	361	−78	1.61	197
O01	1730	428	−49	1.46	284
N01	1686	510	−22	1.35	380
D01	1769	616	3	1.28	488
J02	1521	709	21	1.17	619
F02	1504	808	37	1.13	730
M02	1478	899	48	1.15	895
A02	1480	981	54	1.21	1091
M02	1726	1070	61	1.30	1291
J02	1759	1143	64	1.43	1595
J02	2137	1223	67	1.58	1877
A02	2436	1310	71	1.67	2113
S02	2425	1393	73	1.63	2227
O02	2355	1482	76	1.48	2136
N02	2499	1588	82	1.38	2097
D02	2442	1694	87	1.30	2140
J03	2069	1780	87	1.17	2087
F03	1992	1856	85	1.12	2108
M03	1958	1918	80	1.13	2227
A03	1990	1963	73	1.18	2409
M03	2222	2003	67	1.27	2651
J03	2525	2039	60	1.40	2958
J03	2789	2066	54	1.55	3327
A03	3017	2088	47	1.64	3542
S03	3232	2120	44	1.62	3485
O03	3198	2165	44	1.48	3195
N03	3028	2207	44	1.38	3048
D03	2985	2255	45	1.31	2938
J04					2692

of forecasting until at least one year of the seasonal pattern is available. Thus, exponential smoothing with trend and seasonality needs at least one complete year of data before it is "warmed up" and can start to forecast fairly effectively.

To provide a forecast for any period more than one in the future (April 2004, for example), it is merely a task of taking the most recent value of level that has been calculated (in this case, December 2003), adding to it the most recent value of trend that has been calculated (also in this case, December 2003), times the number of months into the future we wish to forecast (because April is four months past December, it would be times four), and multiplying this value by the seasonal adjustment for April of last year (2003). For April 2004, the calculations are:

$$F_{A04} = (L_{D03} + (T_{D03} \times m)) \times SA_{A03}$$

$$F_{A04} = (2255 + (45 \times 4)) \times 1.18$$

$$F_{A04} = 2874$$

Now that we have introduced the components of trend and seasonality into our basic exponential smoothing formula, we can return to the idea of how to set the value of the smoothing constants. However, now it is not simply a matter of choosing a value for α, but one of choosing values for β and γ, as well. In fact, the accuracy of exponential smoothing with trend and seasonality is very sensitive to the values chosen for the smoothing constants, so this is no small matter.

Adaptive Exponential Smoothing With Trend and Seasonality

As with regular adaptive smoothing, there are several techniques that are adaptive and consider trend and seasonality. One of the most complex computationally is called the Self Adaptive Forecasting Technique (SAFT) and was developed more than 35 years ago (Roberts & Reed, 1969). SAFT is a heuristic technique that examines different combinations of α, β, and γ to arrive at the most accurate forecast. For each forecast each period, SAFT tries each combination of α, β, and γ starting with a value of 0.05 for each and incrementally increasing the values by 0.05 until a value of 0.95 for each is reached. For each of these

6,859 ($19 \times 19 \times 19$, where the 19 is the number of values between 0 and 1, incrementing by 0.05 at a time) combinations, SAFT starts at the beginning of the time series and forecasts using exponential smoothing with trend and seasonality, and it records the resultant value of MAPE. Once the lowest MAPE value combination of α, β, and γ is determined, a local search for a lower MAPE is implemented by examining the values of α, β, and γ above and below each value (including the original three values) at a rate of change of 0.01.

For example, if the first search found the lowest value of MAPE to come from the combination of $\alpha = 0.15$, $\beta = 0.20$, and $\gamma = 0.30$, SAFT would then try all the combinations of α = 0.11, 0.12, 0.13, 0.14, 0.15, 0.16, 0.17, 0.18, 0.19; β = 0.16, 0.17, 0.18, 0.19, 0.20, 0.21, 0.22, 0.23, 0.24; and γ = 0.26, 0.27, 0.28, 0.29, 0.30, 0.31, 0.32, 0.33, 0.34. These 729 ($9 \times 9 \times 9$) combinations are compared to the original best MAPE combination and, again, the lowest combination is chosen.

It should be clear by now that SAFT is a very computationally cumbersome technique (after all, it requires 7,588 trial forecasts for each product each period before it actually makes a forecast) and, as a result, is in little use today. More computationally efficient versions of SAFT try to calculate values of α, β, or γ and use a heuristic similar to SAFT for the smoothing constants that are not directly calculated.

As with adaptive smoothing, because these adaptive exponential smoothing techniques with trend and seasonality essentially all work equally well, we will only discuss the simplest of this group of techniques here. This adaptive smoothing approach, called Adaptive Extended Exponential Smoothing (AEES), uses the absolute value of the percent error from the previous period's forecast to adjust the value of α for the next period's forecast and uses the SAFT heuristic to adjust the values of β and γ (Mentzer, 1988). Thus, the exponential smoothing with trend and seasonality formulae (Equations [9], [10], [11], and [12]) are still used,

$$L_t = \alpha\,(S_t/SA_{t-C}) + (1-\alpha)\,(L_{t-1} + T_{t-1})$$

$$T_t = \beta\,(L_t - L_{t-1}) + (1-\beta)\,T_{t-1}$$

$$SA_t = \gamma\,(S_t/L_t) + (1-\gamma)\,(SA_{t-C})$$

$$F_{t+m} = (L_t + (T_t \times m)) \times SA_{t-C+m}$$

but after each period's sales are recorded, the value of α is adjusted for the next period by Equation (5), repeated here as:

$$\alpha_{t+2} = |\ (F_{t+1} - S_{t+1})/S_{t+1}\ | = |PE_{t+1}|$$

Because this calculation can still produce values outside the range of α, this calculation is again adjusted by the following rules:

If $|PE_{t+1}|$ is equal to or greater than 1.0, then $\alpha_{t+2} = 0.99999$

If $|PE_{t+1}|$ is equal to 0.0, then $\alpha_{t+2} = 0.00001$

Once the new value of α has been calculated, AEES tries each combination of β and γ starting with a value of 0.05 for each and incrementally increasing the values by 0.05 until a value of 0.95 for each is reached. For each of these 361 (19 × 19, where the 19 is the number of values between 0 and 1, incrementing by 0.05 at a time) combinations, AEES starts at the beginning of the time series and forecasts using exponential smoothing with trend and seasonality and records the resultant value of MAPE. Once the lowest MAPE value combination of the calculated value of α and the heuristic values of β and γ is determined, a local search for a lower MAPE is implemented by examining the values of β and γ above and below each value (including the original two values) at a rate of change of 0.01.

For example, if the first search found the lowest value of MAPE to come from the calculated value of α = 0.15, and combination of β = 0.20, and γ = 0.30, AEES would then try all the combinations of β = 0.16, 0.17, 0.18, 0.19, 0.20, 0.21, 0.22, 0.23, 0.24; and γ = 0.26, 0.27, 0.28, 0.29, 0.30, 0.31, 0.32, 0.33, 0.34. These 81 (9 × 9) combinations are compared to the original best MAPE combination and, again, the lowest combination is chosen.

It should be clear by now that AEES is a much less computationally cumbersome technique than SAFT. AEES requires 442 trial forecasts for each product each period before it actually makes a forecast, rather than the 7,588 trial forecasts of SAFT. Further, the exact value of α is calculated rather than the approximation obtained from the SAFT heuristic.

We can illustrate the adaptability of AEES by forecasting the times series with trend and seasonality used in the last section. (See Figure 3.17.) Notice that the forecast in Figure 3.17 "tracks" the demand better

Figure 3.17 AEES With Trend and Seasonality as a Forecast

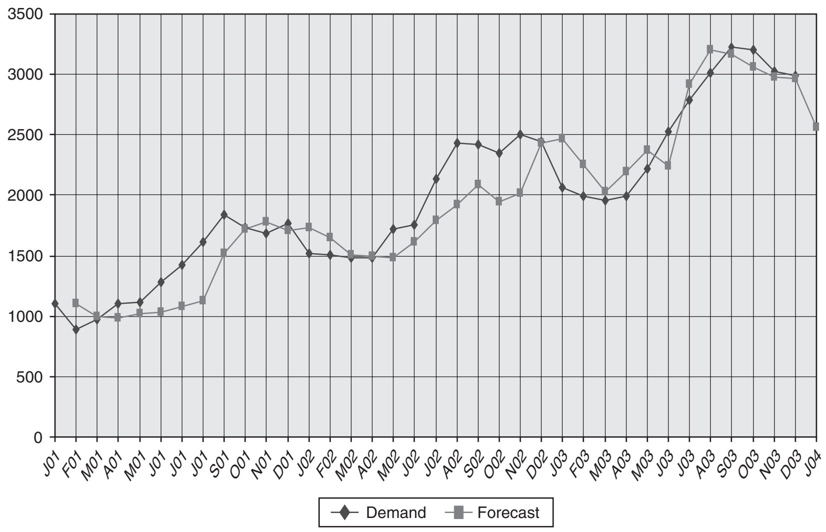

than the forecast in Figure 3.16. This is due to the adaptability of α, β, and γ. Also, the year-to-date MAPE for exponential smoothing with trend and seasonality (Figure 3.16) at the end is 9.73%, while the same calculation for AEES (Figure 3.17) is 7.00%.

❖ FIXED-MODEL TIMES SERIES TECHNIQUES SUMMARY

Considerable effort has been devoted over time to testing the various FMTS techniques discussed here (and variations on these techniques as well) over a wide variety of time series and forecasting horizons and intervals. (For a summary of these efforts, see Mentzer & Gomes, 1994.) To date, no FMTS technique has shown itself to be clearly superior to any of the other FMTS techniques across a wide variety of forecasting levels and time horizons. For this reason, it is recommended that FMTS users keep in mind where the general category of techniques works well and the time series scenario for which each technique was designed.

In general, FMTS techniques should be used when a limited amount of data is available on anything other than an actual history

Table 3.4 FMTS Technique Selection Guidelines

Time Series Component Characteristics	*FMTS Technique*
Stable Level, with No Trend or Seasonality	Exponential Smoothing
Changing Level, with No Trend or Seasonality	Adaptive Smoothing
Level and Trend	Exponential Smoothing with Trend
Level, Trend, and Seasonality	Exponential Smoothing with Trend and Seasonality
Changing Level, Trend, and Seasonal Patterns	AEES

of sales (that is, little data on outside factors such as price changes, economic activity, promotional programs, and so on). This lack of outside (exogenous) data precludes the use of regression (discussed in the next chapter). Further, FMTS techniques are useful when the time series components change fairly regularly. That is, the trend rate changes, the seasonal pattern changes, or the overall level of demand changes. FMTS is much more effective at adjusting to these changes in time series components than are the OMTS techniques to be discussed next, which require more data with stable time series components over a long period of time.

In terms of which FMTS to use in which situations, a general guideline is provided in Table 3.4. However, remember that these are only general guidelines, and it is best to incorporate these techniques into a system (such as the one discussed in Chapter 6) that allows the system to try each FMTS technique on each forecast to be made and select the one that works best in terms of accuracy.

With these general guidelines established, we will now move on to a discussion of the open-model time series (OMTS) techniques.

❖ OPEN-MODEL TIMES SERIES TECHNIQUES

Open-model time series (OMTS) techniques assume that the same components exist in any time series—level, trend, seasonality, and noise—but take a different approach to forecasting these components. Where FMTS techniques assume that certain components exist in the time series and use one set of formulae to forecast this series (that is, the formulae are "fixed"), OMTS techniques first analyze the components in the time series to see which exist and what is their nature. From this information, a set of forecasting formulae unique to that time series is built (that is, the formulae are "open" until the time series components are analyzed).

Various forms of OMTS exist, including decomposition analysis (Shiskin 1961a, 1961b), spectral analysis (Nelson, 1973), fourier analysis (Bloomfield, 1976), and auto-regressive moving average (ARMA) or Box-Jenkins analysis (Box & Jenkins, 1970). All of these OMTS techniques have in common the fact that they first try to analyze the time series to determine the components and, as a result, require a considerable amount of history before any forecasts can be made. For instance, many OMTS techniques recommend no less than 48 periods of data prior to using the technique. Obviously, this is a disadvantage for situations where a limited amount of history is available.

OMTS techniques also have in common the need for considerable understanding of quantitative methods to properly use the techniques. The analysis with OMTS can become quite complex and require considerable input from the forecaster. For these reasons (large data requirements and considerable user experience), OMTS techniques have seen limited use in practice (Mentzer & Kahn, 1995; Mentzer & Cox, 1984a). Improvements in systems technology have made OMTS techniques easier to use (as we will see in Chapter 6), but the data requirements still limit their use.

As with FMTS techniques, there is no evidence that the performance of one of these OMTS techniques is clearly superior to any of the others. Thus, we will again only discuss the simplest of the OMTS techniques here. This technique is called decomposition analysis. To demonstrate decomposition analysis, we will use the time series presented at the beginning of the chapter in Figure 3.1.

Like all OMTS techniques, the purpose of decomposition analysis is to decompose the data into its time series components. The first step in doing this is to remove noise and seasonality from the original time series. As we discussed earlier in the chapter, one of the characteristics of a moving average is that it dampens out any noise and dampens out any regular pattern of fluctuation that has a pattern length that is equal to the number of periods in the moving average. Thus, one of the first things we have to do in decomposition analysis is make a judgment about how long the seasonal pattern is.

Visual examination of Figure 3.1 will, we hope, lead us to conclude that the seasonal pattern takes 12 months. Therefore, a 12-month moving average should remove noise and seasonality from the time series. As in the discussion earlier in the chapter, the value of the moving average in any given period is our estimate of level, and how much that level estimate changes from one period to the next is our estimate of trend. However, because our purpose here is not to forecast, but to decompose the data, we will perform this moving average calculation in a slightly different way than previously discussed. This calculation is as follows:

$$MA_t = (S_{t-5} + S_{t-4} + S_{t-3} + S_{t-2} + S_{t-1} + S_t + S_{t+1} + S_{t+2} + S_{t+3} + S_{t+4} + S_{t+5} + S_{t+6})\ /\ 12$$

Notice that this is a *centered moving average,* which means that we take an average of 12 months and assign that value to the month in the center. The purpose of this is to find a more accurate estimate of the level. If we placed the moving average value at the end of the 12 months used in the calculation, it would have too much old data (lower trend) to accurately represent the level for that period. Conversely, if we place the moving average value at the beginning of the 12 months used in the calculation, it would have too much new data (higher trend) to actually represent the level at that period. Thus, the best place to position this estimate of level is in the center of the periods used in its calculation.

Because this moving average contains the level and the trend, we can simply take the difference between each period to determine the trend. Similarly, since the moving average contains the level and the trend, if we subtract it from the original time series (which contained level, trend, seasonality, and noise), the result is a series of data that contains only the seasonality and the noise. These calculations are demonstrated in Table 3.5.

We now have decomposed the original time series into the level and the trend. All that is left is to remove the noise from the data series containing seasonality and noise, and we will have our final component, seasonality. Again, to remove noise we use an average. However, because each month of the year represents a different season, we want to perform this average calculation within each season. Thus, we take all the January values and average them, then take an average of all the February values, and so on for all 12 months. This calculation is shown

Table 3.5 Decomposition Analysis

Month	*Demand*	*Level and Trend*	*Trend*	*Seasonality and Noise*	*Seasonality*	*Forecast*
J01	1104					
F01	885					
M01	976					
A01	1101					
M01	1120					
J01	1276	1376		−100	−140	
J01	1419	1411	35	8	60	
A01	1615	1463	52	152	261	
S01	1836	1505	42	331	325	
O01	1730	1536	32	194	200	
N01	1686	1587	51	99	204	
D01	1769	1627	40	142	165	
J02	1521	1687	60	−166	−203	
F02	1504	1755	68	−251	−308	
M02	1478	1804	49	−326	−396	
A02	1480	1856	52	−376	−440	
M02	1726	1924	68	−198	−258	
J02	1759	1980	56	−221	−140	
J02	2137	2026	46	111	60	
A02	2436	2067	41	370	261	
S02	2425	2107	40	319	325	
O02	2355	2149	43	206	200	
N02	2499	2190	41	309	204	
D02	2442	2254	64	188	165	
J03	2069	2309	54	−240	−203	
F03	1992	2357	48	−365	−308	
M03	1958	2424	67	−466	−396	
A03	1990	2494	70	−504	−440	
M03	2222	2539	44	−317	−258	
J03	2525	2584	45	−59	−140	
J03	2789		45		60	
A03	3017		90		261	
S03	3232		135		325	
O03	3198		180		200	
N03	3028		225		204	
D03	2985		270		165	
J04			315		−203	2696

Table 3.6 Decomposition of Seasonality

Month	*2001*	*2002*	*2003*	*Average*
January		−166	−240	−203
February		−251	−365	−308
March		−326	−466	−396
April		−376	−504	−440
May		−198	−317	−258
June	−100	−221	−59	−140
July	8	111		60
August	152	370		261
September	331	319		325
October	194	206		200
November	99	309		204
December	142	188		165

in Table 3.6, and the resultant values are added to Table 3.5 in the Seasonality column. Notice that this is not a *multiplicative seasonal adjustment* like we used in FMTS. Rather, it is an *additive seasonal adjustment;* to determine the seasonal adjustment, we add it to (not multiply it by) the level plus trend.

We now have our most recent estimate of level (2,584 in June 2003), our most recent estimate of trend (45 units per month from June 2003), and our most recent estimates of the additive seasonal adjustments for the last 12 months. To forecast a future period (such as January 2004), we take the last estimate of level and add to it the trend times the number of periods into the future. To this value, we add the seasonal adjustment. For January 2004, the calculation is:

$$\begin{aligned}\text{Forecast}_{\text{Jan04}} &= \text{Level}_{\text{June03}} + (7 \times \text{Trend}_{\text{June03}}) + \text{Seasonality}_{\text{Jan03}} \\ &= 2584 + (7 \times 45) - 203 = 2696\end{aligned}$$

This example illustrates just how much data are required to complete OMTS analysis. Although we have 3 years of monthly data in this example, for all but June, only two values were available to estimate the seasonality adjustment for each season (month). With another year's data, three values would be available for each season, which should improve the seasonality adjustment estimates. However, one of the primary drawbacks to OMTS is this dependency on a large amount of data.

❖ SUMMARY

In this chapter, we have covered a number of time series techniques. All have in common a recognition of the time series components—level, trend, seasonality, noise. FMTS techniques deal with these components by assuming certain components are (and are not) in the data, while OMTS techniques analyze the data to determine which components exist. This greater level of sophistication in OMTS is somewhat ameliorated by the considerable data requirements for analysis.

Another characteristic of all the techniques included in this chapter is the fact that they ignore other factors that might have influenced demand, such as price changes, advertising, trade promotions, sales programs, competitive actions, economic activity, and so on. In many cases, much of what time series techniques classify as noise can be explained by looking at these "exogenous" factors. In the next chapter, we turn our attention to regression analysis, a technique that considers these exogenous factors.